

Quantum Monte Carlo for minimum energy structures

Lucas K. Wagner and Jeffrey C. Grossman
*Department of Materials Science and Engineering,
Massachusetts Institute of Technology,
Cambridge, MA 02139
E-mail: lkwagner@mit.edu, jcg@mit.edu
(Dated: May 4, 2010)*

We present an efficient method to find minimum energy structures using energy estimates from accurate quantum Monte Carlo calculations. This method involves a stochastic process formed from the stochastic energy estimates from Monte Carlo that can be averaged to find precise structural minima while using inexpensive calculations with moderate statistical uncertainty. We demonstrate the applicability of the algorithm by minimizing the energy of the $\text{H}_2\text{O}-\text{OH}^-$ complex and showing that the structural minima from quantum Monte Carlo calculations affect the qualitative behavior of the potential energy surface substantially.

First principles electronic structure methods have been used to describe and explain a wide range of properties for different condensed matter systems. A critical step is the accurate determination of the ground state atomic structure, since many important properties of a material can change dramatically depending on the structure. Because of the balance between accuracy and computational cost, density functional theory (DFT) has become a commonly used method to find equilibrium geometries of both molecules and extended systems. The primary reason for this is the availability of forces with little extra computational cost over the energy calculation. Using a typical quasi-Newton minimization algorithm, the local minimum of a potential energy surface can be found in $\mathcal{O}(N_{\text{DOF}})$, where N_{DOF} is the number of degrees of freedom to be optimized. This favorable scaling has made it possible to find minima for many systems of interest. However, in many situations, including transition metals, excited states, and weak binding, current density functional theories may not be accurate enough even for structures, and more accurate post-Hartree-Fock methods that scale as $\mathcal{O}(N_e^{5-7})$, where N_e is the number of electrons, can often be too computationally expensive.

Quantum Monte Carlo (QMC), a stochastic approach to solving the many-body Schrödinger equation, offers favorable scaling to obtain the total energy, $\mathcal{O}(N_e^{2-3})$, and has been shown to provide near chemical accuracy in many different systems[1, 2]. However, there are two major challenges in using QMC methods to obtain high precision minimum energy structures. The first is that the techniques so far proposed to calculate forces in diffusion Monte Carlo all have large variances and error terms that depend on the quality of the trial wave function, which is often poor in systems where DFT fails and one would like to apply QMC methods. In fact, despite much work in recent years[3–8], QMC forces using the highly accurate diffusion Monte Carlo method have not been applied to more than a few degrees of freedom, although the simpler variational Monte Carlo technique has been applied to more[9]. The second challenge is the stochastic nature of Monte Carlo algorithms, which provides uncertainty in any estimator that only decreases as the square

root of the computer time. Reducing the uncertainty enough to resolve the minimum structure accurately using forces or total energies is often prohibitively expensive computationally. Methods such as the stochastic gradient approximation[10, 11] that are able to operate in the presence of noise suffer from this large uncertainty in the forces. As a result, there are no geometry optimizations of more than three[12] degrees of freedom, to our knowledge.

In this article, we describe an algorithm that uses the already-accurate total energies from QMC to obtain minimum energy geometries with well-defined stochastic uncertainties with multiple degrees of freedom. The algorithm consists of two major parts. One is a sequence of minimizations along one dimensional lines. The use of 1D minimizations allows us to use efficient fits to determine the minimum precisely. The second part is a quadratic fit of the many-dimensional energy surface to determine the new search directions. Both of these parts are completely aware of the stochastic uncertainty present in the evaluations of the energy, an important feature obtained by the use of Bayesian inference. We apply this approach to the hydrogen-transfer model of $\text{H}_2\text{O}-\text{OH}^-$ and show that our method can help clarify challenging problems that require accurate calculation of the electronic ground state.

DFT and Hartree-Fock calculations were performed using the GAMESS[13] package. We used soft pseudopotentials[14] to remove the core electrons and a TZP-quality contracted gaussian basis to represent the one-particle orbitals. All-electron calculations were performed using the aug-cc-pVQZ[15] basis to check the basis set and pseudopotential errors. All QMC calculations were performed using the QWalk[16] program. For energies, we used fixed node diffusion Monte Carlo (DMC) with a time step of 0.02-0.05 Hartrees⁻¹, converged for the properties of interest. The trial function was a Slater-Jastrow function with hybrid PBE0[17] orbitals, one and two body terms in the Jastrow, and further checks of the localization error with a three-body Jastrow factor. Further details can be found in e.g. Refs [16, 18].

Our minimization algorithm is similar in spirit to many

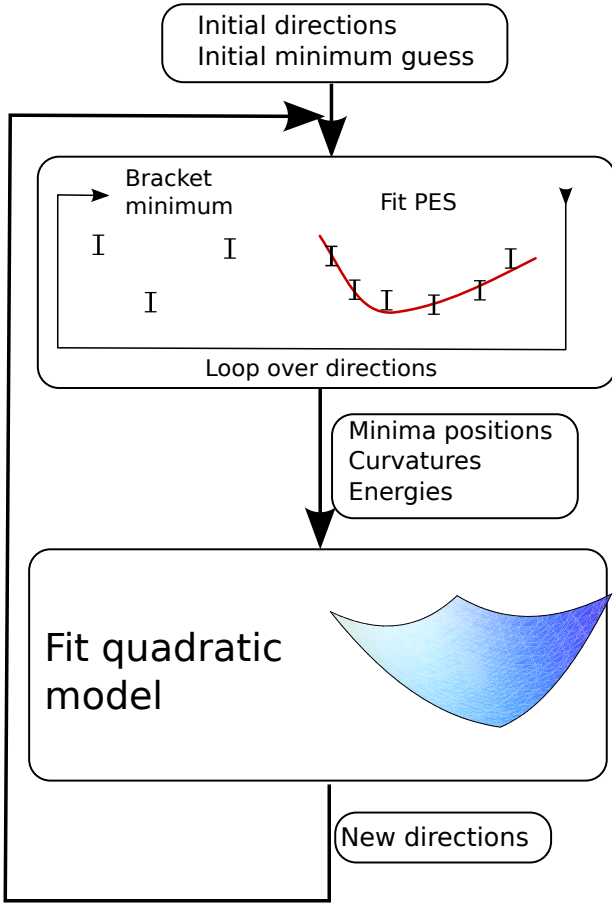


FIG. 1: (color online) An outline of the algorithm. See text for details.

successful minimization algorithms in that it is based on minimizing along directions and updating an estimate of the Hessian matrix to find the diagonal directions. However, in our approach the Hessian matrix is inferred using the Bayesian interpretation of probability, which has the effect of making the algorithm very robust to stochastic uncertainty. Here we present the case when only the value can be calculated, but gradients of the objective function can be included easily if they are available, increasing the efficiency. By using inference, we are able to make very efficient use of the available data to find the Hessian matrix without having to reduce the stochastic uncertainties to small values that cost large amounts of computational time to achieve.

The kernel of the algorithm is a sequence of line minimizations. We show that a sequence of uncertain line minimizations will obtain the true minimum on average. Such a sequence can then be viewed as a generator of random variables whose expectation value is the true minimum. Suppose on the first step, we start with the approximate minimum $\mathbf{x}^{(0)}$. We then define $\delta\mathbf{x}^{(0)} = \mathbf{x}^{(0)} - \mathbf{m}$, where \mathbf{m} is the vector of the unknown true minima. We wish to design our process such that

the expectation value of $\delta\mathbf{x}^{(n)}$ equals zero as $n \rightarrow \infty$. Within the quadratic region around the minimum, the potential energy surface is given by

$$E(\mathbf{x}) = \delta\mathbf{x}^T H \delta\mathbf{x} + E_0, \quad (1)$$

where H is the symmetric Hessian matrix and E_0 is the minimum total energy. On minimizing along each direction i , there are two components of the distance from the true minimum. The first is deterministic and comes about from non-zero δx_j for $j \neq i$. The second is the stochastic error from the uncertainty in the line minimization, which we can estimate using the Bayesian techniques above. We find that

$$\delta x_i^{(1)} = \chi_i^{(1)} + \sum_{j \neq i} \frac{H_{ij}}{H_{ii}} \delta x_j^{(0)}, \quad (2)$$

where $\chi_i^{(1)}$ is a random number. For the n th iteration,

$$\delta x_i^{(n)} = \chi_i^{(n)} - \sum_{j \neq i} \frac{H_{ij}}{H_{ii}} \chi_j^{(n-1)} + \sum_{j \neq i} \sum_{k \neq j} \frac{H_{ij} H_{jk}}{H_{ii} H_{jj}} \chi_k^{(n-2)} + \dots \quad (3)$$

The minimum along each line is found using line fitting, as described in the EPAPS document. This allows for a very efficient determination of the minimum.

One can see from Eqn 3 that the smaller the off-diagonal matrix elements of the Hessian are, the less interference directions have on each other—for a diagonal Hessian, only one minimization in each direction is necessary once we are in the quadratic regime. We can use the information from the line minimizations to estimate the Hessian as the minimization proceeds. We parameterize the quadratic region with a set of parameters \mathbf{c}_Q , including the elements of the Hessian matrix, the minima, and the minimum energy. After performing the line fits, we have distributions of line fit parameters, each given by \mathbf{c}_ℓ for line ℓ . Using Bayes' theorem, the likelihood function of the quadratic parameters L is given by

$$L(\mathbf{c}_Q | D) \propto p(D | \mathbf{c}_Q) = \prod_\ell \int p(D_\ell | L_\ell) p(L_\ell | \mathbf{c}_Q) d\mathbf{c}_\ell, \quad (4)$$

where D is the set of function evaluations and stochastic uncertainties given by e.g. QMC and each D_ℓ is the subset of function evaluations used to minimize along a given line ℓ . Since $p(L_\ell | \mathbf{c}_Q)$ is not based on stochastic data, it is a delta function that forces \mathbf{c}_ℓ to be consistent with \mathbf{c}_Q as follows. Since in the quadratic region, $E(\mathbf{x}) = \delta\mathbf{x}^T H \delta\mathbf{x} + E_0$, minimizing along a direction \mathbf{v} from a starting position \mathbf{x}_0 gives the one-dimensional function of t , the position along the line:

$$E(t) = (\mathbf{x}_0 + t\mathbf{v} - \mathbf{m})^T H (\mathbf{x}_0 + t\mathbf{v} - \mathbf{m}) + E_0. \quad (5)$$

This gives the following constraints on each set of line parameters for the minimum, curvature, and minimum function value c_ℓ :

$$t_m = - \frac{\mathbf{v}^T H (\mathbf{x}_0 - \mathbf{m}) + (\mathbf{x}_0 - \mathbf{m})^T H \mathbf{v}}{2\mathbf{v}^T H \mathbf{v}} \quad (6)$$

$$\left. \frac{d^2 E}{dt^2} \right|_{t_m} = 2\mathbf{v}^T H \mathbf{v} \quad (7)$$

$$E(t_m) = (\mathbf{x}_0 + t_m \mathbf{v} - \mathbf{m})^T H (\mathbf{x}_0 + t_m \mathbf{v} - \mathbf{m}) \quad (8)$$

The net result of this transformation is a set of parameters that includes the Hessian matrix, the location of the minimum, the objective function at the minimum, and the parameters used for the line fits above the three constraints implied by the \mathbf{c}_Q 's. We examine the properties of this probability distribution function in the EPAPS material. It turns out that the maximum likelihood estimator is typically accurate enough to determine the Hessian, which is then diagonalized to obtain new search directions.

We outline a single iteration of the algorithm in Fig 1. This step operates in two principle regimes. The first is far away from the minimum. In this regime, the Hessian inference method behaves similarly to a deterministic direction set method, with the directions being determined by the Hessian inference. The second regime is when the calculation has mostly converged and the deterministic error in Eqn 2 is small compared to the stochastic error. In this regime, deterministic direction choices are particularly useless, but the inferred Hessian method is able to automatically account for the stochastic error. Once in this “stochastic regime,” we can use Eqn 3 to justify averaging the $\mathbf{x}^{(n)}$'s to obtain a more precise estimate of the minimum. For stochastic functions, the performance of the Hessian inference is much higher than traditional methods such as Powell's method[19], which we compare in Fig 2. Since Powell's method uses the concept of conjugate directions to find the search directions, the uncertain minimizations cause the algorithm to fail quickly. Even when the directions are reset every sweep, they can become linearly dependent within a single sweep, causing a high failure rate for even moderate dimensionality.

To show the value of optimizing geometries within QMC, we apply the method to the $\text{H}_2\text{O}-\text{OH}^-$ complex (Fig 3), which is present in liquid water and important in many systems in condensed matter, biology, and chemistry. The shape of this potential energy surface is crucial to understanding hydrogen transfer in water. In this case, we use our knowledge of the system to choose fixed search directions for efficiency reasons, omitting the Hessian inference method. For systems that are not as easily decomposed, the Hessian inference is invaluable. It has been noted[20] that current DFT functionals disagree on the ground state structure of this complex. The potential minima are the non-centrosymmetric structure (A) and the centrosymmetric structure (B).

Hartree-Fock and second order Møller-Plesset perturbation theory (MP2) find that structure A is lower in energy, with a barrier to transfer the proton. This is the traditional picture of this structure. The local density approximation and generalized gradient approximation (PBE[21]) of DFT find that structure B is lower in energy. Using our QMC line minimization method without




FIG. 2: (color online) Estimates of reliability for the Hessian inference method versus Powell line minimization. There are several runs for each number of degrees of freedom, each with a different potential energy surface and relative starting position. The potential surface was a randomly generated positive definite quadratic surface with a random minimum and a small cubic perturbation. The condition number of the matrices was typically 50 or less. We added a Gaussian random variable to the values of the potential energy surface to simulate uncertain evaluation. Convergence was judged when the RMS distance to the true minimum converged to the stochastic floor of the minimization.

TABLE I: $\text{H}_2\text{O}'\text{-OH}^-$. The non-QMC methods are optimized using conjugate gradient routines in GAMESS. The asterisked DMC result corresponds to constraining the DMC geometry search to the symmetry of structure B.

Method	O-O'	O'-H	Structure type
LDA	2.448	1.224	B
PBE	2.470	1.235	B
MP2	2.469	1.123	A
DMC	2.491(2)	1.111(3)	A
DMC*	2.469(3)	1.235(2)	B

constraints (seven degrees of freedom) we find structure A to be the minimum energy. The results are summarized in Table I. DMC differs *qualitatively* from the DFT results in that structure A is the minimum, and quantitatively with MP2, since the oxygen-oxygen distance in MP2 is much smaller than in DMC for structure A.

In Fig 3, we present the DMC energies of several minima. Without the geometry optimization algorithm, we would use minima from some other methods, for example, PBE and MP2 to obtain structures B and A respectively, then evaluate the total energy using DMC. In this case, since the MP2 approximation obtains a poor O-O distance, DMC predicts a much higher energy for structure A and thus favors the centrosymmetric geometry. How-




FIG. 3: (color online) The relative DMC energies of geometries obtained by minimizing different potential energy surfaces. The line is a guide to the eye. Stochastic error bars are the size of the symbols.

ever, this is qualitatively incorrect, as we can see from the DMC-optimized structures, which predict structure B to be about 0.015 eV higher in energy than structure A. It is notable that 0.015 eV is a very small energy difference on the scale of chemical bonding, which is why a small error in the geometry by using a lower level theory is enough to reverse the ordering. This energy difference is an upper bound to the barrier to transfer the hydrogen, and thus at room temperature the hydrogen is free to transfer between the oxygen atoms. This may have important implications for the development of effective solvent models for water.

In this work, we have viewed the DMC potential energy surface as a given without attempting to improve the accuracy over a simple Slater-Jastrow form; however, our method can also be used to optimize the nodal

surface directly and thus further improve the accuracy of the DMC calculation. Similarly, any Monte Carlo method that depends on continuous parameters could put the stochastic line minimization algorithm to use. Stochastic line minimization may find application in experiments whose objective is to obtain a precise minimum in many-dimensional space, but the experiments are difficult to perform precisely. If it is instead easier to perform many iterations of the experiment, then the minimization method that we have outlined may be applicable. The scaling for a quadratic potential is $\mathcal{O}(N_{DOF}^2)$ if only values are used; if even approximate gradients are available, the scaling can be reduced to approximately $\mathcal{O}(N_{DOF})$, as discussed in the EPAPS document.

The algorithm is quite general, requiring only evaluations of the objective function with a known distribution, and can thus be applied to many minimization problems. At its core, it is a method for converting noisy value calculations into precise minima with rigorous error bounds, so it can be applied in any of the many problems where that is necessary. In electronic structure, diffusion Monte Carlo is accurate and applicable to many systems beyond the ground state of molecules, including solid structures and excited states[22]. The scaling for DMC geometry optimization is then $\mathcal{O}(N_{DOF}^{1-2}N_e^{2-3})$, depending on the size of the system and whether gradients are available. This favorable scaling combined with the already known high accuracy of quantum Monte Carlo could open up new levels of accuracy in structure determination in condensed matter and chemical physics.

This work was supported by the Department of Energy under grant No. de-sc0002623 and the NSF under grant No. 0425914. We wish to thank NERSC and Teragrid as well for providing computational resources. We would also like to acknowledge Yosuke Kanai for useful discussions.

-
- [1] J. Kolorenč and L. Mitas, *Physical Review Letters* **101**, 185502 (pages 4) (2008), URL <http://link.aps.org/abstract/PRL/v101/e185502>.
 - [2] J. C. Grossman, *The Journal of Chemical Physics* **117**, 1434 (2002), URL <http://link.aip.org/link/?JCP/117/1434/1>.
 - [3] S. Chiesa, D. M. Ceperley, and S. Zhang, *Physical Review Letters* **94**, 036404 (pages 4) (2005), URL <http://link.aps.org/abstract/PRL/v94/e036404>.
 - [4] R. Assaraf and M. Caffarel, *J. Chem. Phys.* **119**, 10536 (2003).
 - [5] M. Casalegno, M. Mella, and A. Rappe, *J. Chem. Phys.* **118**, 7193 (2003).
 - [6] A. Badinski and R. J. Needs, *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)* **76**, 036707 (pages 10) (2007), URL <http://link.aps.org/abstract/PRE/v76/e036707>.
 - [7] A. Badinski, J. R. Trail, and R. J. Needs, *The Journal of Chemical Physics* **129**, 224101 (pages 7) (2008), URL <http://link.aip.org/link/?JCP/129/224101/1>.
 - [8] C. Filippi and C. J. Umrigar, *Phys. Rev. B* **61**, R16291 (2000).
 - [9] C. Attaccalite and S. Sorella, *Physical Review Letters* **100**, 114501 (pages 4) (2008), URL <http://link.aps.org/abstract/PRL/v100/e114501>.
 - [10] H. Robbins and S. Monro, *Ann. Math. Stat.* **22**, 400 (1951).
 - [11] A. Harju, B. Barbiellini, S. Siljamäki, R. M. Nieminen, and G. Ortiz, *Phys. Rev. Lett.* **79**, 1173 (1997).
 - [12] C. A. Schuetz, M. Frenklach, A. C. Kollias, and J. W. A. Lester, *The Journal of Chemical Physics* **119**, 9386 (2003), URL <http://link.aip.org/link/?JCP/119/9386/1>.
 - [13] M. W. Schmidt, K. K. Baldridge, J. A. Boatz, S. T. Elbert, M. S. Gordon, J. H. Jensen, S. Koseki, N. Matsunaga, K. A. Nguyen, S. J. Su, et al., *J. Comput Chem*

- 14, 1347 (1993).
- [14] M. Burkatzki, C. Filippi, and M. Dolg, The Journal of Chemical Physics **126**, 234105 (pages 8) (2007).
 - [15] J. Thom H. Dunning, The Journal of Chemical Physics **90**, 1007 (1989), URL <http://link.aip.org/link/?JCP/90/1007/1>.
 - [16] L. K. Wagner, M. Bajdich, and L. Mitas, Journal of Computational Physics **228**, 3390 (2009), ISSN 0021-9991, URL <http://www.sciencedirect.com/science/article/B6WHY-4VGPWS6-3/2/df02c999649199aa85eb26ed66d87850>.
 - [17] C. Adamo and V. Barone, The Journal of Chemical Physics **110**, 6158 (1999), URL <http://link.aip.org/link/?JCP/110/6158/1>.
 - [18] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Rev Mod Phys **73**, 1 (2001).
 - [19] R. Brent, *Algorithms for Minimization without Derivatives* (Prentice-Hall, 1973), pp. 116–155.
 - [20] V. N. Staroverov, G. E. Scuseria, J. Tao, and J. P. Perdew, The Journal of Chemical Physics **119**, 12129 (2003), URL <http://link.aip.org/link/?JCP/119/12129/1>.
 - [21] J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**, 3865 (1996).
 - [22] F. Schautz, F. Buda, and C. Filippi, The Journal of Chemical Physics **121**, 5836 (2004), URL <http://link.aip.org/link/?JCP/121/5836/1>.

Appendix A: Details of the minimization algorithm

In this EPAPS document, we step through the minimization process in detail. Note that some of the details are not necessarily perfectly optimal, but they do work robustly and efficiently. We thus provide this description only as a reference implementation that could potentially be improved significantly.

1. Line minimization

Suppose we wish to minimize a function along a line using the values. For a deterministic function, a good method is the golden ratio method. However, for a non-deterministic function, locating the minimum by comparing the value is bound to be inefficient, since the objective function $f(x)$ scales as $f(x) \propto x^2$, and the work to resolve energy differences scales as the inverse square root of the difference. Therefore, near the minimum, the effort to resolve position differences scales as the inverse *quartic* root of the difference. The computational cost using fitting (and gradients) on the other hand, scales as the inverse square root. Whether fitting using the value of the function or the gradient of the function is more accurate will depend on their relative variances and the specific problem. In this work, we use exclusively values since they are easily attainable in quantum Monte Carlo methods.

We developed a reliable strategy to bracket the minimum. Given a starting position (t_0) and a trust radius r_t ,

the algorithm works with three points, initialized as follows: $t_1 = t_0 - r_t$, $t_2 = t_0$, $t_3 = t_0 + r_t$. Let $f_i = f(t_i) + \chi$ be the function evaluated at t_i with a normal random error χ with standard deviation σ , which can be controlled by performing longer calculations. Define $S(f_i, f_j)$ as being true when f_i is significantly different from f_j , i.e., about four standard deviations away. A basic algorithm as follows:

1. Evaluate f_1 , f_2 , and f_3 at uncertainty σ
2. If $S(f_1, f_2)$ and $S(f_2, f_3)$:
 - If $f_2 < f_3$ and $f_2 < f_1$ stop. We have bracketed the minimum.
 - If $f_2 < f_3$ and $f_2 > f_1$, then set $t_i = t_i - r_t$ for $i = 1, 2, 3$, return to 1.
 - If $f_2 > f_3$ and $f_2 < f_1$, then set $t_i = t_i + r_t$ for $i = 1, 2, 3$, return to 1.
3. Else $\sigma = \sigma/2$ and return to 1.

This algorithm occasionally gets stuck when two points are roughly the same distance from the minimum and thus have nearly the same value. This can be ameliorated by searching between points when only two are insignificantly different.

Once the minimum is bracketed to within the trust region, we can then fill in points along the line and fit. The distribution of points is not very critical, so long as there are a few points far away from the minimum to determine the curvature and higher order parameters accurately. We typically use a four parameter cubic form: $f(x) = c_0 + c_2(x - c_1)^2 + c_3(x - c_1)^3$, which gives a good balance between being accurate, easy to fit, and few parameters. A quadratic form requires trust radii that are very small, reducing the efficiency of the algorithm. We perform the fit using Bayesian inference given the set of data $D = \{t_i, f_i, \sigma_i\}$. The likelihood function is given by

$$L(\mathbf{c}|D) \propto p(D|\mathbf{c}), \quad (\text{A1})$$

where

$$p(D|\mathbf{c}) \propto \prod_i \exp[-(f(t_i, \mathbf{c}) - f_i)^2 / 2\sigma_i^2] \quad (\text{A2})$$

is the probability that we would have obtained the data given a set of parameters. The distribution of a given parameter c_i can be found by taking the marginal distribution of $L(\mathbf{c}|D)$, or for more speed at the cost of not knowing the uncertainty well, simply using the maximum likelihood estimator.

2. Hessian estimation

One begins by noting that near the minimum, the energy is described by a quadratic function:

$$E(\mathbf{x}) = \delta \mathbf{x}^T H \delta \mathbf{x} + E_0. \quad (\text{A3})$$

Far away from the minimum, this is an approximation to speed up the minimum search. Near the minimum, we used this property to prove the average properties of the stochastic sequence of line minimizations. Recall that only the minimum of the line minimization must be in the quadratic region. With no further assumptions, we can estimate the Hessian (and full quadratic surface) given a collection of line minimizations, even in the presence of stochastic noise.

Let \mathbf{c}_Q be the collection of all parameters for the quadratic region, \mathbf{c}_ℓ the collection of line parameters for line ℓ . We wish to find $p(\mathbf{c}_Q|D)$, where D is the set of function evaluations, arranged in lines. By Bayes' theorem, we know that

$$p(\mathbf{c}_Q|D) \propto p(D|\mathbf{c}_Q) = \prod_{\ell} \int p(D_{\ell}|\mathbf{c}_Q)p(\mathbf{c}_\ell|\mathbf{c}_Q)d\mathbf{c}_\ell \quad (\text{A4})$$

Since $p(D_{\ell}|\mathbf{c}_Q)$ is not based on stochastic data, it is a delta function that forces \mathbf{c}_ℓ to be consistent with \mathbf{c}_Q as follows. In the quadratic region, $E(\mathbf{x}) = \delta\mathbf{x}^T H \delta\mathbf{x} + E_0$, minimizing along a direction \mathbf{v} from a starting position \mathbf{x}_0 gives the one-dimensional function

$$E(t) = (\mathbf{x}_0 + t\mathbf{v} - \mathbf{m})^T H (\mathbf{x}_0 + t\mathbf{v} - \mathbf{m}) + E_0. \quad (\text{A5})$$

This gives the following constraints on each set of line parameters \mathbf{c}_ℓ :

$$t_m = -\frac{\mathbf{v}^T H (\mathbf{x}_0 - \mathbf{m}) + (\mathbf{x}_0 - \mathbf{m})^T H \mathbf{v}}{2\mathbf{v}^T H \mathbf{v}} \quad (\text{A6})$$

$$\left. \frac{d^2 E}{dt^2} \right|_{t_m} = 2\mathbf{v}^T H \mathbf{v} \quad (\text{A7})$$

$$E(t_m) = (\mathbf{x}_0 + t_m \mathbf{v} - \mathbf{m})^T H (\mathbf{x}_0 + t_m \mathbf{v} - \mathbf{m}) \quad (\text{A8})$$

We thus only need to integrate over the variables in each line that are not specified by the above equations. We can replace the \mathbf{c}_ℓ by the reduced line parameters $\tilde{\mathbf{c}}_\ell$. Since there are three constraints per line, the dimensionality of $\tilde{\mathbf{c}}_\ell$ is three less than the dimensionality of \mathbf{c}_ℓ . We can recover the full set of line parameters as a function of the quadratic parameters and the reduced line parameters; i.e., $\mathbf{c}_\ell(\mathbf{c}_Q, \tilde{\mathbf{c}}_\ell)$

We eventually want the likelihood function for the quadratic parameters and the reduced line parameters: $L(\mathbf{c}_Q, \{\tilde{\mathbf{c}}_\ell\}|D)$, where D is the data. This is given by

$$L(\mathbf{c}_Q, \{\tilde{\mathbf{c}}_\ell\}|\{D_\ell\}) \propto \prod_{\ell} L(\mathbf{c}_\ell(\mathbf{c}_Q, \tilde{\mathbf{c}}_\ell)|D_\ell). \quad (\text{A9})$$

where $L(\mathbf{c}_\ell|D_\ell)$ is given by Eqn A1.

The first estimate of the Hessian matrix is achieved by minimizing the violation of Eqn A7 from the line fits. The minimum of the quadratic model is estimated as the last minimum found in a line minimization, and the

minimum energy similarly. The parameters are then optimized for maximum likelihood using a conjugate gradient optimization method. Finally, the parameters are randomized and reoptimized to get out of local maxima. We find that the maximum likelihood estimator is typically quite sufficient for determining search directions; it is not necessary to sample to find distributions.

Note that each line minimization places three constraints on \mathbf{c}_Q and \mathbf{c}_Q has a dimensionality of $1 + N_{DOF} + N_{DOF}(N_{DOF} + 1)/2$, so to make this a fully determined problem, $1/6 + N_{DOF}/3 + N_{DOF}(N_{DOF} + 1)/6$ line minimizations are needed. This is always smaller than the N_{DOF}^2 line minimizations necessary in Powell's method, the standard in gradient-free minimization, with the additional benefit that the Hessian inference method is uncertainty-aware, which allows much higher efficiency when uncertainty can be traded for computational time.

If gradient information is available, then they can be integrated directly into the likelihood function from Eqn A9 as a simple product:

$$L(\mathbf{c}_Q, \{\tilde{\mathbf{c}}_\ell\}|\{D_\ell\}, \{\mathbf{g}_i\}) \propto \left[\prod_{\ell} L(\mathbf{c}_\ell(\mathbf{c}_Q, \tilde{\mathbf{c}}_\ell)|D_\ell) \right] \left[\prod_i L(\mathbf{c}_Q|\mathbf{g}_i) \right] \quad (\text{A10})$$

where the likelihood function can be approximated as

$$L(\mathbf{c}_Q|\mathbf{g}_i) \propto \exp[-(\nabla_{\mathbf{x}} E(\mathbf{x}_i, \mathbf{c}_Q) - \mathbf{g}_i)^T C_i (\nabla_{\mathbf{x}} E(\mathbf{x}_i, \mathbf{c}_Q) - \mathbf{g}_i)], \quad (\text{A11})$$

and C_i is the covariance matrix between all the gradient components. The Hessian inference then becomes essentially a quasi-Newton algorithm that is uncertainty aware. Approximate gradients are also useful, since minima can be determined by using only the value to perform the line fits. In some instances, one can imagine elaborations of this where DFT or other approximations to the gradient are used to accelerate the minimum search.

3. A detailed look as the algorithm proceeds

1. Start a line minimization sweep, for each direction
 - Bracket the minimum
 - Fill in data points within the trust region from the minimum
 - Fit the data points with a function
 - Update the minimum
 - Add fit to list of fits
2. Prune list of fits to contain only the last $N_{history}$ sweeps
3. Fit N-dimensional quadratic model to lines, find maximum likelihood
4. Diagonalize Hessian




FIG. 4: Typical convergence of the eigenvalues and the estimate of the minimum for a nine-dimensional optimization problem. The potential energy surface was a random quadratic function with a cubic perturbation. The potential energy surface was calculated as $E(\mathbf{x}) + \chi$, where χ was a random Gaussian variable with mean zero.

5. Update directions

6. Return to 1.

In Fig 4, we show the convergence of the eigenvalues from the Hessian estimation to the exact ones, and the minima from the line minimizations.

Appendix B: Proofs of the properties of the stochastic sequence

We have the following equations:

$$\begin{aligned} \delta x_i^{(n)} = & \chi_i^{(n)} - \sum_{j \neq i} \frac{H_{ij}}{H_{ii}} \chi_j^{(n-1)} \\ & + \sum_{j \neq i} \sum_{k \neq j} \frac{H_{ij} H_{jk}}{H_{ii} H_{jj}} \chi_k^{(n-2)} + \dots \end{aligned} \quad (\text{B1})$$

$$\begin{aligned} \text{Var}(\delta x_i^{(n)}) = & \text{Var}(\chi_i^{(n)}) + \sum_{j \neq i} \frac{H_{ij}^2}{H_{ii}^2} \text{Var}(\chi_j^{(n-1)}) \\ & + \sum_k \text{Var}(\chi_k^{(n-2)}) \left[\sum_j (1 - \delta_{ij})(1 - \delta_{jk}) \frac{H_{ij} H_{jk}}{H_{ii} H_{jj}} \right]^2 \\ & + \dots \end{aligned} \quad (\text{B2})$$

Suppose that $\sum_{j \neq i} |H_{ij}/H_{ii}| \leq x, \forall i$ for some $0 < x < 1$. Then $|\sum_{k \neq j} \frac{H_{ij} H_{jk}}{H_{ii} H_{jj}}| \leq \left| \frac{H_{ij}}{H_{ii}} \right| x$, $|\sum_{j \neq i} \sum_{k \neq j} \frac{H_{ij} H_{jk}}{H_{ii} H_{jj}}| \leq x^2$, and so on.

1. The average is equal to the true minimum

Suppose that $\langle \chi_i^{(j)} \rangle = 0 \forall i, j$. Then

$$\langle \delta x_i^{(n)} \rangle \leq x^n \delta x_i^{(0)}. \quad (\text{B3})$$

This goes to zero exponentially in n , so we can obtain an average arbitrarily close to zero by starting the average at high enough n .

2. The average has finite variance

Assume $\text{Var}(\chi_i) \leq v_m \forall i$. Then

$$\text{Var}(\delta x_i^{(n)}) \leq v_m (1 + x^2 + x^4 + \dots) = \frac{v_m}{1 - x^2}, \quad (\text{B4})$$

for large n .

3. The sequence has a finite correlation time

Suppose that

$$\langle \chi_i^{(k)} \chi_i^{(j)} \rangle = \sigma^2 \delta_{jk} \forall i, j, k \quad (\text{B5})$$

for some variance σ^2 . Then the correlation between elements n and j in the sequence is

$$|\langle \delta x_i^{(j)} \delta x_i^{(n)} \rangle| \leq \sigma^2 \sum_{k=0}^j x^{n-k} = \sigma^2 x^{n-j} \sum_{k=0}^j x^k, \quad (\text{B6})$$

where we reversed the sum indices. As n and j become large while keeping the difference finite, we can then find that

$$|\langle \delta x_i^{(j)} \delta x_i^{(n)} \rangle| \leq \frac{\sigma^2 x^{n-j}}{1 - x}. \quad (\text{B7})$$

This goes to zero exponentially with $n - j$, and since the variance is finite, the correlation coefficient also goes to zero exponentially.